

UNITED STATES PATENT APPLICATION

For

**MULTI-PROCESSOR COMPUTING SYSTEM THAT EMPLOYS
COMPRESSED CACHE LINES' WORTH OF INFORMATION AND
PROCESSOR CAPABLE OF USE IN SAID SYSTEM**

Inventors:

Chris J. Newburn
Ram Huggahalli
Herbert Hum
Ali-Reza Adl-Tabatabai
Anwar M. Ghuloum

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(408) 720-8300

MULTI-PROCESSOR COMPUTING SYSTEM THAT EMPLOYS COMPRESSED CACHE LINES' WORTH OF INFORMATION AND PROCESSOR CAPABLE OF USE IN SAID SYSTEM

Field of Invention

[0001] The field of invention relates generally to computing systems; and, more particularly, to a multi-processor computing system that employs compressed cache lines' worth of information and processors that are capable of use in the multi-processor computing system.

Background

[0002] **Figure 1** shows a portion of an architecture for a basic computing system that includes: 1) a processor 101; 2) a cache 102; 3) a memory controller 103; and, 4) a system memory 104. The processor 101 implements software routines by executing instructions that perform various operations on elements of data. The instructions and data elements are stored in the cache 102 and/or system memory 104. When the processor 101 needs a specific instruction or data element it looks to the cache 102 for the desired instruction or data element before requesting it from system memory 104.

[0003] Generally, cache 102 is deemed to be "faster" than the system memory 104. Better said, the processor 101 waits less time waiting for an instruction or data element that resides in the cache 102 than an instruction or data element that resides in the system memory 104. This disparity in waiting time as between

the cache 102 and system memory 104 typically arises as a consequence of the cache 102 being implemented with inherently faster memory cells (e.g., SRAM cells) than those of which the system memory is implemented (e.g., DRAM cells).

[0004] Per bit of storage space an SRAM type cache 102 is more expensive than a DRAM type system memory 104. The computing system architecture of **Figure 1** therefore attempts to optimize both cost and performance by being designed to store more frequently used instructions and data elements in the cache 102 and less frequently used instructions and data elements in the system memory 104. By storing the more frequently used instructions and data elements in the cache, the processor should endure acceptable “timing penalty hits” in the form of wasted time waiting for instructions/data to be fetched from system memory 104 because a significant percentage of the instructions/data needed by the processor will be found in the cache 102.

[0005] In order to enhance the percentage of “cache hits” (i.e., the instances where a needed instruction or data element is found the cache 102), notions of “temporal locality” and “spatial locality” come into play. Temporal locality is the notion that a single instruction or data element is apt to be used soon after it has already been used. Spatial locality is the notion that instructions and data elements that are located near each other in memory (i.e., have similar addresses) tend to be used at about the same time. Temporal locality is accounted for by keeping instructions and data elements in cache 102 for at least some period of time after they are first transferred from system memory 104 into cache 102.

[0006] Spatial locality is accounted for by designing the cache 102 to be loaded with a block of data from system memory 102 (i.e., multiple instructions or data elements) whose content is proximate to (e.g., “surrounds”) any single instruction or data element that needs to be fetched from system memory 104. For example, if an instruction at address X is needed from system memory 104, instead of transferring only the needed instruction from system memory 104, instead of transferring only the needed instruction from system memory 104 to cache 102, a block of content corresponding to a plurality of addresses that are related to address X is transferred from system memory 104 to cache 102.

[0007] **Figure 2** attempts to depict such a situation by showing that a first contiguous “block” of content 105 (which is referenced through multiple system memory addresses) is loaded into a single cache line 107; and, that a second contiguous “block” of content 106 (which is referenced through a different set of multiple system memory addresses) is loaded into another single cache line 108. For simplicity, **Figure 2** shows the cache 204 as a single structure. Various computing systems are designed with different levels of cache, however. For example, many types of computing systems have two levels of caches (a level one (L1) cache and a level two (L2) cache) where the first level cache (L1) corresponds to less processor waiting time than the second level cache (L2). The L1 cache is supposed to store the most frequently used data elements and instructions while the L2 cache is supposed to store data elements and instructions that are used less frequently than those in L1 cache but more frequently than those in system memory.

[0008] Traditionally, both cache levels are implemented with a faster memory type as compared to system memory (e.g., both L1 and L2 cache are implemented with SRAM memory cells); however, the L1 cache is integrated onto the same semiconductor die as the processor while the L2 cache is implemented with different semiconductor die than the processor. As “on chip” cache accesses are faster than “off chip” cache accesses, accesses to the L1 cache correspond to less waiting time for the processor than accesses to the L2 cache.

[0009] The memory controller 103 is responsible for taking requests from the processor 101 for data, that are not satisfied by the cache, and managing the process of servicing those requests in system memory 104. There may be many different kinds of requests, such as load requests for data that is not present in the cache, and evictions of data from the cache that need to be stored back into memory. Typically, the memory controller is able to pipeline requests, so that many requests may be outstanding, and can be serviced in parallel with a much shorter average latency. The memory controller is responsible for interfacing with the details of a particular memory technology, and isolates the system memory from the processor in a modular fashion. The memory controller may either be integrated with the processor, e.g. on the same die, or may be separated, e.g. in a chipset.

[0010] The system memory is typically implemented with a specific type of system memory (e.g., EDO RAM, SDRAM, DDR, etc.).

Figures

[0011] The present invention is illustrated by way of example, and not limitation, in the figures of the accompanying drawings in which like references indicate similar elements and in which:

[0012] **Figure 1** shows a portion of a computing system's architecture;

[0013] **Figure 2** shows that a block of content from system memory is loaded into a single cache line;

[0014] **Figure 3a** shows an improved approach in which aligned blocks of system memory content can be compressed into a single cache line;

[0015] **Figure 3b** shows an exemplary resulting map of substantive system memory content after compressed cache lines are evicted from cache and stored into system memory;

[0016] **Figure 4a** shows a portion of a single processor architecture capable of using compressed cache lines;

[0017] **Figure 4b** shows a compression map that can be utilized by a system memory controller to keep track of those blocks within its system memory that have been compressed;

[0018] **Figure 4c** shows a memory controller having a compression map cache and compression/decompression logic;

[0019] **Figure 4d** shows a memory controller having register space for identifying physical system memory allocation for a compression map;

[0020] **Figure 5a** shows a pair of methods that can be utilized by the memory controller of **Figure 4a** during a write of a cache line into system memory;

[0021] **Figure 5b** shows a trio of methods that can be utilized by a memory controller during a read of a cache line from system memory;

[0022] **Figure 5c** shows a trio of methods related to referring to the contents of a compression map;

[0023] **Figure 6a** shows a first embodiment of a memory controller;

[0024] **Figure 6b** shows a second embodiment of a memory controller;

[0025] **Figure 6c** shows a third embodiment of a memory controller;

[0026] **Figure 7** shows a multiprocessor computing system that is capable of using compressed cache lines' worth of information;

[0027] **Figures 8a through 8c** show various transfers involving compression that may occur within a multi-processor computing system, such as the multi-processor computing system of **Figure 7**, that is capable of working with compressed cache line's worth of information;

[0028] **Figure 9** shows an embodiment of a design for a data-path within a compression/decompression hub;

[0029] **Figure 10** shows an embodiment of a map of the address bits used to reference a cache line's worth of information from a cache that is capable of storing both compressed and uncompressed cache lines' worth of information;

[0030] **Figure 11** shows an embodiment of a first set of cache coherency rules;

[0031] **Figure 12** shows an embodiment of a second set of cache coherency rules;

- [0032] **Figure 13** shows an embodiment of a third set of cache coherency rules;
- [0033] **Figure 14** shows an embodiment of a fourth set of cache coherency rules;
- [0034] **Figure 15** shows an embodiment of a methodology that involves a request for ownership of a compressed cache line's worth of information;
- [0035] **Figure 16a** illustrates a traditional memory address implemented in a cache;
- [0036] **Figure 16b** illustrates one embodiment of a memory address implemented in a cache capable of compressing/decompressing cache lines' worth of information;
- [0037] **Figure 17** illustrates one embodiment of a tag array entry for a cache capable of compressing/decompressing cache lines' worth of information;
- [0038] **Figure 18** is a block diagram illustrating one embodiment of a cache controller;
- [0039] **Figure 19** illustrates one embodiment of a set and way selection mechanism in a cache capable of compressing/decompressing cache lines' worth of information;
- [0040] **Figure 20** illustrates one embodiment of byte selection logic.

Detailed Description

Compression of Cache Lines' Worth of Information

[0041] As a matter of clarification, a cache line is a basic unit of storage space in a cache. For example, in many applications a unique tag and set address are used to specially identify a single cache line within a computing system's cache. A cache line is therefore implemented with specific electrical circuitry elements (e.g., SRAM cells). By contrast, "a cache line's worth of information" or "line of information" is an amount of information (e.g., data elements or instructions) that can fill a cache line. Here, recalling the discussion of **Figure 2**, the amount of information stored at "block" 105 corresponds to a cache line's worth of information because the content of block 105 fills cache line 107.

[0042] **Figure 3a** demonstrates an approach that expands upon the notion of spatial locality so as to compress, into a single cache line 307, a pair of aligned system memory 304 blocks 305, 309 that would ordinarily occupy a pair of cache lines (i.e., the information of a pair of cache lines' worth of information are compressed into a single cache line). A second instance is also observed in **Figure 3a** in which the content of another pair of aligned memory blocks 306, 310 that would ordinarily occupy a pair of cache lines are compressed together so as to occupy a second single cache line 308.

Compression is a technique that reduces the amount of data needed to express information (such as an instruction or a data element) without impacting the substantive content of the message itself (i.e., without eliminating the ability to

recapture the “number” used to represent the instruction or data element). The ability to compress a pair of aligned blocks into a single cache line should result in faster computing system performance because the effective size of the cache is increased (; and, therefore, the likelihood of needing to incur the longer access latency to a slower cache level or system memory is decreased). Moreover, as described in more detail below, computing system bandwidth can be enhanced by suppressing access to information because it is compressed with other information that already been accessed.

[0043] In an embodiment, referring to **Figures 3a** and **Figure 4a**, the processor’s cache controller 410 is fitted with compression/decompression logic 411 that compresses two cache lines worth of information together if: 1) the cache lines worth of information represent aligned, contiguous blocks of memory; and, 2) the informational content of the pair of the cache lines worth of information is capable of compression into a single cache line. The type of compression employed may take on various forms such as Lempel-Ziv, Wilson-Kaplan, X-Match or perhaps other known or proprietary types of compression.

[0044] In an embodiment, to say that companion blocks are aligned means that the lowest address associated with the companion blocks is a multiple of the combined size of the companion blocks. For example, if each cache line’s worth of information is 64 bytes, then the base address, N , of two contiguous cache lines worth of information (i.e., a first at N referred to as the “lower” cache line worth of information and a second at $N+64$ referred to as the “higher” or “upper” cache line worth of information) is divisible by 128 (i.e. the remainder of $N/128$ is

0). As a further example, referring to **Figure 3a**, block 305 would be addressable with a base address of N; and, block 309 would be addressable with an address of $N + 64$.

[0045] For convenience, aligned contiguous cache lines worth of information are referred to as “companions” of one another. Thus, in light of the preceding paragraph, a pair of companion cache lines’ worth of information are compressed by the compression/decompression logic 411 if their substantive content is capable of compression. Likewise, the compression/decompression logic 411 is capable of decompressing a compressed cache line worth of information into two separate companion cache lines’ worth of information if a write occurs to the content of a cache line that causes the content to no longer be compressible into a single cache line.

[0046] Once a pair of cache lines’ worth of information have been compressed together they are treated by the entire computing system as a single cache’s line worth of information (e.g., by being referenced with a single tag and set address while being stored within a cache) until a write occurs to the compressed information that causes it to be deemed no longer compressible. Upon being deemed non compressible, the information is split into a pair of companion cache lines. Thus, if the processor 401 continually accesses from cache 402 a specific cache line’s worth of compressed information, the processor 401 continually reads only a single cache line even though the informational equivalent of a pair of cache lines are actually being read. Toward the end of this detailed description are details regarding exemplary compression/decompression logic

implementation(s) that may be instantiated, for example, in a processor whose cache is capable of compressing/decompressing cache lines' worth of information.

[0047] Once information has been compressed into a single cache line's worth of information, the single cache line's worth of information may be treated as any "normal" uncompressed cache line of information such as: 1) being read/written from/to a cache (including a particular level of cache); 2) being read/written from/to a system memory; and, 3) any structure designed to transport a cache line's worth of information (such as, to name a few: a front side bus or point-to-point link that transports cache lines' of information between a processor and a memory controller that controls access to a system memory; and, in a multiprocessor environment, a pair of processors that share cached information).

[0048] As an example of a possible transfer of a cache line's worth of compressed information, referring to **Figure 4a**, consider a situation where a cache line's worth of information is evicted from the cache 402 (e.g., because it has not been used frequently enough to be deemed worthy of continued storage in the cache 402) and transferred to system memory 404. Here, the cache line's worth of compressed information can be stored in the system memory addressing space of a single block that can store a single cache line's worth of information. For example **Figure 3b** shows a depiction of the utilization of system memory 304, with respect to the substantive content of blocks 305 and 309 of **Figure 3a**, after the compressed content of cache line 307 has been

evicted from the cache 302, 402. **Figure 3b** shows that upon eviction from cache 302, 402 and storage into system memory 304, 404 the content of cache line 307 is stored so as to occupy only memory block 305. This is in stark contrast to the utilization of the system memory that existed prior to compression, shown in **Figure 3a**, for storing the same amount of information.

[0049] Note that even though two cache lines' worth of data may be stored in the space normally occupied by a single cache line's worth of data, when stored in system memory in compacted form, this does not imply an effective increase in system physical memory capacity, as is true for compressed caches. This is because in system memory, the address space is not compacted. Compacting the address space requires modifications to page tables and thus it requires operating systems support, which the schemes presented here are capable of avoiding entirely.

[0050] That is, after compression only block 305 is needed to store the informational content of that which was earlier stored in blocks 305 and 309 prior to compression. **Figure 3b** also demonstrates that, upon eviction, the compressed contents of cache line 308 are restored in system memory 304 so as to only occupy block 306 even though blocks 306 and 310 were used to store the same information prior to compression. If one of the "compressed content" blocks 305, 306 of **Figure 3b** is needed again by the processor 401, it is read from system memory 304, 404 by memory controller 403a as a single cache line's worth of information and is transferred (again as a single cache line's worth

of information) from memory controller 403a to processor 401 and written (again as a single cache line's worth of information) into the processor's cache 402.

Memory Controller

[0051] In the context of single processor environments the memory controller may behave largely without any recognition or cognizance of the compression/decompression activity taking place. That is, for example, the processor 401 may "keep track of" and manipulate those cache lines worth of information that are compressed and those that are not compressed; and, by contrast, the memory controller is designed to simply read and write blocks of data in accordance with identifiers or labels assigned by the processor 401.

[0052] However, a more sophisticated memory controller 403a that takes into account which blocks of system memory are used to store content that corresponds to compressed cache lines worth of information (and/or which blocks of system memory are used to store content that corresponds to non-compressed cache lines worth of information) may be able to reduce the demand for system memory accesses so as to make the system memory's usage more efficient within the computing system. For example, by refusing to read a second block of data because its substantive content has just been read from a compressed, first block of data, the demand that is exercised on the system memory is effectively reduced.

[0053] As a more detailed example, consider a multiprocessor environment where the processors are capable of compressing information into its cache

lines. Here, a first processor (e.g., processor 401 in Fig. 4a) may compress information into a cache line and then subsequently evict it from its cache 402 so that it is stored into system memory 404. If a second processor in the multi-processor system (not shown in **Figure 4a**), without knowledge of the first processor's compression activity, desires to read from system memory 404 information stored in both companions of the compressed information, the memory controller 403a may be designed to be "smart enough" to only read the compressed cache line's worth of information in response to receiving a pair of read requests from the second processor (i.e., a first request for the first companion and a second request for the second companion). Here, the compressed cache line's worth of information will be sufficient to satisfy both requests made by the second processor.

Compression Map

[0054] **Figure 4b** provides a trio of embodiments 412a, 412b, 412c for a body of information, referred to as a compression map 412, that may be used by the memory controller 403a to recognize the existence of compressed information within its system memory 404. Firstly, referring to "basic embodiment" 412a, note that the compression map 412a may be stored as a bit map in system memory 404 that identifies, for each block of information in system memory 404, whether that block's corresponding cache line's worth of information is currently stored in system memory 404 in a compressed format or in a non-compressed format.

[0055] In typical implementation, an address column is not actually included in the compression map (e.g., in cases where the map covers the whole memory).

Figure 4b shows an address column in each of embodiments 412a, 412b, 412c so that the reader can easily understand a compression map's organization and structure. Specifically, bits have been provided an active value "1" (while others have not been provided an inactive value "0") in the context of examples that are based upon the system memory shown in **Figure 3b** and that are discussed immediately below. As such, the compression map may be implemented as a data structure being organized to have specific values at locations (e.g., data fields) that correspond to specific system memory blocks.

[0056] Compression map embodiment 412a of **Figure 4b** is depicted so as to apply to the system memory observed in **Figure 3b**. Specifically, recall that the system memory of **Figure 3b** stores information in block 305 that corresponds to the compression of information that existed in blocks 305 and 309 prior to compression. Because the information of blocks 305 and 309 of **Figure 3a** have been compressed together, the compression map 412a of **Figure 4a** provides an indication (a "1") for each of these blocks 305, 309. Likewise, because the information of blocks 306 and 310 of **Figure 3a** have been compressed together (into block 306), the compression map 412a of **Figure 4b** provides an indication for each of these blocks as well 305, 309. Referring to **Figure 4a**, note that the compression map 412a may be stored in the system memory itself 404.

[0057] A "more elaborate" compression map embodiment 412b of **Figure 4b** includes bitmap information as described above with respect to embodiment

412a as well as additional information in the form of: 1) information (e.g., in select cases such as instance 414) that provides the substantive content of a cache line's worth of information; 2) indication(s) 415 of the type of compression used for each cache line's worth of information that is stored in a compressed format.

[0058] The former additional information 414 corresponds to an extreme form of compression that may be applied: a) to the content of system memory blocks having non-compressed cache lines' worth of information; and/or b) "on top of" those cache lines' worth of information that are already stored in system memory in a compressed format (embodiment 412b indicates a single instance of the former). For example, if the cache line's worth of information that is stored in a particular system memory block is "all zeroes"; then, a single "zero" (e.g., zero 414) may be stored at the particular block's location in the compression map 412b. Similar indications may be used for any type of constant value (e.g., "all 1s"). Here, the memory controller would be expected to include logic (such as summation logic (e.g., the sum of all zeroes will be zero)) that identifies those cache lines having a constant value.

[0059] The later form of additional information 415 indicates a specific type of compression. Here, recall that different types of compression may be employed (e.g., Lempel-Ziv, Wilson-Kaplan, X-Match, etc.). Not only may compression of only a single type exist within any one particular computing system (e.g., a single system that only uses Lempel-Ziv); but also, embodiments may also be crafted where a single system is capable of implementing different types of compression

(e.g., a single system that can use any of the Lempel-Ziv, Wilson-Kaplan, X-Match and perhaps other compression algorithms).

[0060] Both of the compression map embodiments 412a, 412b show a bit that provides compressed/uncompressed status for each aligned block in system memory that can store a cache line's worth of information. By contrast, embodiment 412c uses only one bit to represent the compressed/uncompressed status of each pair of aligned system memory blocks. Here it is worthy to note that compression ratios other than 2:1 may be employed (such as 4:1); and that, the size of a compression map that is implemented according to the approach of embodiment 412c will become smaller as the compression ratio increases. That is, for 2:1 compression, a bit is used to represent every aligned pair of memory blocks; while, if a 4:1 compression ratio were used, there would be a bit for every group of four aligned memory blocks. Note also that the more elaborate information of embodiment 412b can be added to embodiment 412c.

[0061] Alternate embodiments of the compression map could use selective, hierarchical schemes, rather than a flat bit vector. A flat bit vector must have one bit for every block in memory. Compression may be applied selectively to only certain regions of memory, and thus the compression map could be made to cover only those regions of memory that are subject to compression. Likewise, compression may actually have occurred (so far) in only a subset of memory regions, even though additional regions may be subject to compression. The various sections of the compression map that cover the regions which have been fully or partially compressed can be linked together as a linked list, or worked into

a hierarchy of data structures that cover progressively smaller regions and sub regions of memory.

[0062] Recalling that a condition for the compression of a pair of companions is that the substantive content of the companions “be compressible” into a single cache line’s worth of information, and owing to the different mathematical techniques employed across different compression schemes, a particular compression technique may regard a pair of particular companions to be compressible while other compression schemes may not regard the same companions to be compressible (e.g., the substantive content of a pair of companions may be compressible under Lempel-Ziv but not under Wilson-Kaplan or X-Match). As such, more companions are likely to be compressed in a computing system that “offers” different types of compression as compared to a computing system that offers only a single type of compression. The compression type indication 415 of the enhanced bit map embodiment 412b of **Figure 4b** can be used in such a system (noting that it indicates compression type “A” was used for blocks 305, 309 and compression type “B” was used for blocks 306, 310). Therefore, compression/decompression logic 411 of **Figure 4a** should be understood to be capable of performing singular or multiple types of compression depending on the particular embodiment.

[0063] Also, note from **Figure 4a** that the compression map may be stored in system memory 404. In an embodiment, the memory controller 403a is designed to fetch a portion of the compression map 412 from system memory 412 at an

appropriate moment to check upon the compression/decompression status of one or more system memory blocks.

[0064] In order to reduce the efficiency penalty associated with accessing system memory 404 in order to fetch a portion of the compression map 412, note also that the memory controller 403a is designed to include a compression map cache 413. The compression map cache 413 contains one or more recently fetched portions of the compression map. Similar to a normal cache, compression map information may be continuously updated in the compression map cache 413 until evicted to system memory 404. As described in more detail below with respect to **Figures 5a** through **5d**, the compression map cache 413 is referred to when compression map information is desired. If the desired information is not found in the compression map cache 413, the information is fetched from the compression map 412 that resides in system memory 404.

[0065] **Figure 4c** demonstrates that a memory controller 403b configured to work with a compression map 412 may be instrumented not only in a computing system having a single processor 420; but also, with one or more processors (such as processor 420 and perhaps other processors not shown in **Figure 4c**) that do not possess the ability to compress/decompression their cached information. Thus, the memory controller 403b of **Figure 4c** is capable of being the main (and perhaps only) component in the computing system that is conscious of any compression activity. The depiction of **Figure 4c** therefore, in contrast to **Figure 4a**, shows that the memory controller 403b itself can be retrofitted with the appropriate compression/decompression logic 416 used for

compressing and decompressing cache lines (noting also that processor 420 is devoid of such logic). The compression/decompression logic 416 may support one or more types of compression/decompression techniques.

[0066] The memory controller 403b may further include a compression map cache 413 as described above in reference to **Figure 4a**. In working with processor(s) that do not maintain any cognizance of compression/decompression activity, the memory controller 403b presents/receives uncompressed cache lines worth of data to/from the processor(s). Specific methodologies that may be executed by a memory controller 403b that is operating in an environment where the processor(s) can't operate with compressed cache lines are described in more detail further below.

[0067] **Figure 4d** is meant to convey that the compression map 412 may be stored within a "physical" continuous addressing range of the system memory 404 rather than being implemented in a "virtual" fashion across unrelated memory locations (with, for example, link listing techniques that are managed in software). By implementing the compression map 412 across a physical addressing space, the Operating System (OS) may operate without awareness of the compression/activity; which, in turn, saves the OS from being bogged down with executing instructions for managing or recognizing which locations of system memory 404 are to be used for the compression map 412. As such, a significant degree of overhead is avoided from being imparted upon the OS.

[0068] By configuring the compression map to be implemented across a physical range of the system memory's addressing space, the compression map

should also be capable of being managed and controlled by the computing system's hardware rather than its operating system. As discussed above, this should "free up" the OS so as to be substantially unburdened with overhead relating to the compression map. In an embodiment, the Basic Input Output System (BIOS) 430 indicates what specific physical address range of the system memory 404 is to be used for the compression map 412 by causing a pair of registers 431, 432 to be written into. For example, a first address might be stored into register 431 that defines the starting address of the compression map; and, a second address might be stored into register 432 that defines the ending address of the compression map.

[0069] Alternatively, the size of the compression map might be stored into one of registers 431, 432 while a starting or ending address is stored in the other of registers 431, 432 (noting that the size of the compression map might vary depending on whether 2:1, 4:1 or another compression aspect ratio is employed). Subsequent to the loading of registers 431 and 432 the hardware is capable of refraining from storing non compression map information into the addressing space identified through registers 431 and 432; and, likewise, directing the compression map only toward the same addressing space. The registers 431, 432 may alternatively be located in a processor.

[0070] If the compression map is physically distributed across multiple local memories, or a compression map scheme is used that does not require each portion of the compression map to reside in physical contiguous memory, more

than one pair of registers may be used to communicate from the BIOS to the hardware where the compression map resides.

[0071] It is also worthwhile to note that storing the compression map across a contiguous physical address range that is hidden from, and not paged by, the operating system should permit the compression map to be referenced using physical addresses without having to handle changes in the virtual address and page faults that may occur as the operating system swaps pages out of physical memory and into virtual memory and back again. This is another way in which this scheme avoids the need for OS support, and is transparent to software.

Compression Map Uses

[0072] As discussed above, the compression map represents whether particular cache lines' worth of information stored in main memory are compressed or uncompressed. In various embodiments it is updated with each write to memory that changes the compression state of that memory. A compression map can be used for at least the following three purposes: 1) to effectively change the target address of an upper cache line's worth of information line that has been compressed in a non-duplicative scheme; 2) to decide whether a cache line's worth of information that has just been read from system should be decompressed or not by a memory controller that performs decompression; and, 3) to suppress a system memory access if requests for separate companions are recognized and the companions have been compressed. Each of these are discussed more fully below in the context of writes to system memory and reads from system memory.

System Memory Writes

[0073] **Figure 5a** shows a pair of memory controller methods 551, 552 for writing a cache line's worth of information into a block of system memory. Each of the methods 551, 552 of **Figure 5a** invoke a compression map. According to the first methodology 551, a compressed cache line's worth of information is received by the memory controller (e.g., as sent from a processor) 501. The compressed cache line's worth of information is presumed to be identified to the memory controller as being in a compressed format (e.g. with a set bit in a control header or an activated line). In response to the reception of the compressed cache line's worth of information, the memory controller updates 502 the compression map to reflect that the received cache line's worth of information is compressed. Any of embodiments 412 – 412c of **Figure 4b** or variants thereof can be used to implement the compression map.

[0074] In order to perform the update 502, referring to **Figure 4a**, the memory controller 403a refers to the compression map cache 413. If the section of the compression map that is correlated with the system memory block to which the received compressed cache line's worth of information is associated resides within the compression map cache 413; then, only the compression map cache 413 is updated (so as to avoiding accessing the compression map 412 in system memory 404). If the appropriate portion of the compression map is not within the compression map cache 413, the appropriate portion is fetched from system memory 404 and updated 502.

[0075] Note also that in an embodiment (such as that depicted in **Figur 4c**) where the memory controller 403b is coupled to a processor that does not use cache lines with compressed information, process 501 would be slightly modified such that: 1) only an uncompressed cache line's worth of information would be received at box 501; 2) between boxes 501 and 502 the memory controller 403b would determine that the received cache line's worth of information is compressible with its companion (e.g., by referring to the substantive content of its companion in an inbound or outbound queue of the memory controller 403b); and, 3) prior to execution of box 503 the memory controller 403b would compress the received cache line's worth of information with its companion.

[0076] Recall that two companion cache lines worth of information correspond to a pair of aligned blocks of address space in main memory. Here, the combination of a pair of aligned blocks can be viewed as a larger "macro block" of memory space; where, one companion occupies the "lower half" of the macro block, and the other occupies the "upper half" of the macro block, when they are each uncompressed. When the companions are compressed, the substantive content of the entire macro block can be referenced with the addressing information used for only one of the smaller companion blocks (e.g., the addressing information used for the lower half of the macro block). When uncompressed, the upper and lower halves of the macro block are separately addressable.

[0077] For example, referring briefly back to **Figures 3a** and **3b**, the combination of blocks 305 and 309 can be viewed as a macro block of information where block 305 corresponds to the “lower half” of the macro block (because it is referenced using the lower addressing space of the pair of blocks 305, 309) and block 309 corresponds to the “upper half” of the macro block (because it is referenced using the higher addressing space of the pair blocks 305, 309). When uncompressed, “lower half” 305 is separately addressable and “upper half” 309 is separately addressable. When compressed, the combined content of both halves can be accessed by addressing lower half 305.

[0078] The memory controller should be designed to recognize, for any uncompressed cache line’s worth of information, which half of a macro block it is supposed to occupy and which half of a macro block its corresponding companion is supposed to occupy. For example, referring briefly back to **Figures 3b** and **4b**, the memory controller would be designed to recognize that an uncompressed cache line’s worth of information that is addressed to upper half 309 is the companion line of an uncompressed cache line’s worth of information that is addressed to lower half 305. Such recognition is straightforward based upon the mathematics of the alignment scheme that defines which blocks are companions one another. For simplicity, a lower half of a macro block will hereinafter be referred to as a lower block and a higher half of a macro block will be referred to as a higher block.

[0079] For 2:1 compression ratios, a pair of embodiments are possible as to the usage of the upper and lower blocks of a macro block when its substantive content is compressed. Referring back to **Figure 5a**, in a first embodiment referred to as “non-duplication”, irrespective of whether a compressed cache line of information to be written into system memory was compressed by the memory controller or a processor, the write 503 of a compressed cache line’s worth of information involves a write to the address space of only the lower block of the corresponding macro block. **Figures 3a** and **3b** illustrate a “non-duplication” approach because, as originally discussed, if blocks 305 and 309 of **Figure 3a** are compressed together, only the lower block 305 of **Figure 3b** is written to (of course, alternatively, only the higher block could be written to).

[0080] According to a “non-duplication” approach, as described in more detail below with respect to methodology 555 of **Figure 5b**, the memory controller refers to the compression map prior to a read because a request (e.g., by a system component that is unaware of any compression activity) for a higher block that has been compressed into a lower block can only be satisfied by reading from the lower block (i.e., the target specified in the request is different than the location in system memory from where a read is performed to satisfy the request).

[0081] In an alternative second embodiment, referred to as “duplication”, the write 503 of a compressed cache line involves a write to the address space of all the blocks among the applicable companion set (e.g., both the lower and higher blocks among the applicable companion set for 2:1 compression). For example,

for a 2:1 compression approach, if blocks 305 and 309 of **Figure 3a** are compressed together, both blocks 305 and 309 of **Figure 3b** are written to with the same compressed information. The duplication approach allows the memory controller to avoid having to retrieve information from a lower compressed block of information when a request for the upper block's information is received (as described just above with respect to the "non duplication" embodiment).

[0082] As such, the compression map does not need to be referred to for requests for "upper" blocks of information. If **Figure 3b** were to be modified to reflect a duplicative approach, upper block 309 would be shaded and it would be further understood that the content of upper block 309 is the same compressed content as that stored in lower block 305. Likewise, upper block 310 would be shaded and it would be further understood that the content of upper block 310 is the same compressed content as that stored in lower block 306.

[0083] In the second memory controller write methodology 552 of **Figure 5a** an uncompressed cache line is received 504 from a processor that is capable of performing compression. As such, the received, uncompressed cache line is deemed "uncompressible" for whatever reason. The compression map is therefore updated 505 (e.g., by writing a "0" in the compression map at a location that represents the uncompressed cache line's corresponding block) and written into system memory 506.

[0084] Write methodology 552 could also be slightly modified to represent a write process in systems where the memory controller performs compression/decompression (such as a system as described in **Figure 4c** where

the processor does not support compression). As such, unlike the immediately preceding discussion, it is unknown whether the received uncompressed cache line is compressible or uncompressible. In such a case, between boxes 504 and 505, the compression/decompression logic 416 of the memory controller decides that the received cache line is not compressible (e.g., by analyzing its content along with the content of its companion as found in an input queue or output queue of the memory controller). If it were deemed compressible, it would be compressed with its companion and write 506 would be a write of compressed information.

System Memory Reads

[0085] **Figure 5b** shows a trio of memory controller read methods 553, 554, 555. The first read method embodiment 553 is directed to implementations, such as that depicted in **Figure 4c**, where the memory controller performs the compression and decompression of cache lines and the processor(s) with whom the memory controller communicates do not use cache lines that support compressed information. As such, for any cache line's worth of information that is read 507 from system memory, the memory controller refers 508 to the compression map to see if the information being read is compressed (note that the reference 508 to the compression map is shown as being after the read 507 but may alternatively be performed in parallel with and/or prior to the read 507). If the read cache line's worth of information is compressed the memory controller decompresses it 509, 510. If the read cache line's worth of information is not

decompressed the memory controller does not attempt to decompress it 509, 511.

[0086] If the memory controller happens to reside in a computing system having components that recognize the existence of compressed cache lines' worth of information; then, the memory controller may be implemented without compression/decompression logic (e.g., the environment of **Figure 4a** is applicable rather than the environment of **Figure 4c**). If so, the memory controller should be designed so as to simply signify whether the read information is compressed or decompressed (e.g., by adjusting a value within a header that is appended to the cache line's worth of information) rather than actually perform decompression. To represent a read process for such a memory controller, box 510 of methodology 553 of **Figure 5b** should correspond to providing an indication (e.g., in a header or activated line) that the read information is compressed and box 511 should correspond to providing an indication the read information is not compressed.

[0087] Methodologies 554 and 555 may be performed by a memory controller that has compression/decompression logic or a memory controller that does not have compression/decompression logic. The second read methodology 554, which has already been briefly alluded to, involves the memory controller being designed to be "smart enough" to avoid making a second read to system memory for a companion of an already read compressed cache line's worth of information. According to this methodology, if the memory controller recognizes that there are pending read requests for cache lines' worth of information that are

companions of one another, the compression map is referred to 512, 514. If the compression map reveals that the companions are compressed together the memory controller only reads 518 the compressed cache line from system memory in order to satisfy both requests.

[0088] If the compression map reveals that the companions are not compressed together, the memory controller reads both cache line's worth of information (for a 2:1 compression scheme) separately 516, 517 from their corresponding lower and upper blocks of information in order to satisfy the requests. If there are no pending read requests for cache lines' worth of information that are companions of one another the memory controller behaves like a normal memory controller and simply performs a separate read 513 from system memory to satisfy each request.

[0089] It is worthwhile to note that the term "pending" request means that the physical memory component has not, as yet, actually responded to the memory controller that issued the request. However, it is possible for the memory controller to suppress a second request even if the physical memory component has already responded to the first (i.e., the first request is no longer "pending"). For example, the memory controller could be designed to suppress any second request for compressed information provided the data for the second request can be provided (e.g. from the memory controller) from the results of the first request. Therefore, the ability to suppress requests can be extended to situations beyond those described by methodology 554 of **Figure 5b**.

[0090] In cases where the memory controller is designed to perform decompression, the memory controller may perform both of read methods 553 and 554 together in a continuous flow; where: 1) methodology 554 is largely performed prior to the read, 2) methodology 553 is largely performed after the read, and, 3) any of reads 518, 517, 513 of methodology 554 also correspond to read 507 so as to “connect” methodologies 553, 554 together. If the methodologies 553, 554 are connected in this fashion note that reference 508 may be “skipped” (i.e., not performed) if reference to the compression map 514 was made prior to the memory read. This is so because the answer to inquiry 509 that methodology 553 indicates is to be performed after the read can be gleaned from reference 514 which is made prior to the read.

[0091] Methodology 555 corresponds to a write methodology that can be used to implement the “non-duplication” write approach discussed above with respect to **Figure 5a**. Here, the compression map is referred to if the target address of the requested cache line’s worth of information corresponds to the upper block for the companion pair 519, 521. If the requested cache line’s worth of information has been compressed, the compressed cache line is read from the lower block 522, 520. If the requested cache line’s worth of information has not been compressed, the uncompressed requested cache line’s worth of information is read from the target block specified in the read request. If the target block specified in the read request is not the upper block, the memory controller simply reads a compressed or uncompressed cache line’s worth of

information from the system memory with addressing that corresponds to the lower block 519, 520 (i.e., no reference to the compression map is needed).

[0092] Similar to methodology 554, methodology 555 may be combined with methodology 553 for a memory read performed by a memory controller that also performs decompression. Here, either of reads 523 and 520 of method 555 can be viewed as read 507 of method 553 so as to connect the two methodologies 555, 553 together. If the execution of method 555 flows through compression map reference 521 prior to memory read 507, compression map reference 508 can be skipped because the answer to inquiry 509 can be gleaned from reference 521. Use of a duplicative scheme as discussed with respect to **Figure 5a** removes the need for methodology 555 because no change of target address is effected with a “yes” answer to inquiry 522.

Compression Map Cache Lookup

[0093] **Figure 5c** provides a trio of methodologies 556, 557, 558 that are related to the references to the compression maps 508, 514, 521 that were discussed just above with respect to **Figure 5b**. In particular, methodology 556 shows a more detailed depiction of a process that may be used to implement any of compression map references 508, 514, 521 of **Figure 5b**. Methodology 556 corresponds to a basic cache/system memory read process – albeit applied with the novel features of a memory controller’s compression map cache and a compression map residing in system memory. Better said, in order to refer to the compression map 556, the memory controller first refers 525 to its on-board compression map cache 413.

[0094] As the compression map cache 413 only contains a portion of the entire compression map 412 that is stored in system memory 404, if the information for the cache line is found in the compression map cache 526 (i.e., “a hit”) – the reference to the compression map is complete. If the information for a particular block is not found in the compression map cache (i.e., “a miss”), the information is fetched from the complete compression map that resides in system memory 526, 527 (i.e., a read to system memory is performed).

[0095] A compression map miss corresponds to an efficiency penalty because a system memory read 527 is performed as a consequence. Methodologies 557, 558 correspond to methodologies that may be performed by the memory controller in order to mitigate the timing penalty hit associated with a compression map cache miss for either of the compression map references 514, 521 of **Figure 5b** that precede a corresponding memory read 516 – 518, 520, 523. Both of methodologies 557, 558 apply to a memory controller that performs decompression (e.g., because it works with a processor that does not use compressed cache lines as depicted in **Figure 4c**) and therefore performs methodology 553 of **Figure 5b** for all system memory reads of a cache line.

[0096] Methodology 557 shows that the memory controller may be designed to perform the reference to the compression map cache 531 that occurs prior to a memory read 514, 521 in the process of satisfying a second memory read request in a time period that overlaps with the read 530 of a cache line’s worth of information from system memory to satisfy a first memory read request. That is, performing the pre memory read cache lookup 531 and the memory read 530 of

different requests with some degree of parallelism should help mitigate the timing penalty hit if a cache lookup 531 turns out to be a miss. Here, the degree of temporal overlap (e.g., partial or otherwise) between the memory read and the cache lookup may vary depending on implementation.

[0097] In the particular case of a pipelined memory controller and system memory (so as to be capable of servicing multiple system memory read requests in parallel), the read of a cache line's worth of information 530 to service a first request may continue in parallel with the read of compression map information 532 that is needed if the compression map lookup 531 is a miss. Methodology 557 shows such a situation in both flow chart form (subscript "1" in labels 530, 531, 532) and Gantt chart form (subscript "2" in labels 530, 531, 532).

[0098] Methodology 558 is applicable to the "non-duplicated" embodiment discussed above with respect to **Figure 5a**. It shows that prediction (either "compressed" or "uncompressed") may be used in the case of a cache miss in performing references 514, 521; and, that the subsequent reference to the compression map 508 to check if decompression is needed is used to check the validity of the prediction. According to methodology 558, if the cache lookup results in a miss 532, the state of the requested cache line is predicted to be compressed or uncompressed. In a first embodiment, the state is conservatively predicted to be uncompressed. In another embodiment the recent history of the compression map's content is used as a basis for predicting a compressed state or an uncompressed state. The cache line is then fetched in accordance with the prediction.

[0099] For example, if the requested cache line's worth of information corresponds to an upper block and is predicted to be in a compressed state, a cache line's worth of information is read from the address of the lower block 533. Contrarily, if the predicted state of the cache line's worth of information is uncompressed, a cache line's worth of information is read 533 from the address of the upper block. The appropriate portion of the compression map is then fetched from system memory 534 (because miss 532 indicates that the compression map does not contain information for the applicable cache line's worth of information). The proper compression map information is then checked to see if the prediction was correct 535. If so, the remaining read request processes are performed. In a further embodiment, a compression map cache update for another request may occur after execution of inquiry 532 but before the execution of box 534. If so, box 534 may instead correspond to a "re-look" into the compression map; and, if a hit occurs, a fetch to system memory for compression map information may be eliminated altogether.

Memory Controller Embodiments

[00100] Figures 6a through 6c show various memory controller embodiments 603a, 603b, 603c; where, each memory controller embodiment 603 includes a compression map cache 613a, 613b, 613c. Embodiment 613a does not include any compression or decompression logic circuitry. Embodiment 613b includes decompression circuitry 616b. Embodiment 613c includes compression logic circuitry 616d and decompression logic circuitry 616c. For each of the embodiments 613a, 603b, 603c, the bus/point-to-point link

interface(s) 601 correspond to an interface of the memory controller where: 1) requests for memory reads and memory writes are received; 2) responses to the requests are provided. Because requests may conceivably be received from and responded to over a bus (e.g., a front side multidrop bus); and/or, received from and responded over a point-to-point link (e.g., a first inbound link that receives requests and a second outbound link that sends responses), interface 601 may be interface to a bus and/or point-to-point link.

[00101] The request/response queues 602 of each embodiment 603a, 603b, 603c queue requests in the inbound direction (e.g., in a first, request queue). The scheduler logic circuitry 623 of each embodiment 603a, 603b, 603c schedules the servicing of these requests. The memory request queue 604 of each embodiment 603a, 603b, 603c queues requests that have been scheduled by the scheduler logic circuitry 623. The memory interface 605 of each embodiment 603a, 603b, 603c is responsible for reading/writing information from/to the particular type of memory that the memory controller is coupled to. The request/response queues 602 of each embodiment 603a, 603b, 603c also queue responses to requests in the outbound direction (e.g., in a second, response queue).

[00102] In various embodiments, the updates or references 502, 505, 514, 521 discussed above may be performed by the scheduler logic circuitry (or from some other appropriate location). For each of embodiments 603a, 603b, 603c, input 612 to the compression map cache 613 can be viewed in a first instance as an input that supplies compression map information from the external memory to

the compression map (e.g., in the case of a compression map cache miss).

Moreover, input 612 can be viewed in a second instance as the reference to the compression map information that is performed in association with a read of a cache line's worth of information from system memory.

[00103] Here, recall from the above discussion of methodology 553 of **Figure 5b** that if the memory controller is capable of performing decompression - e.g., embodiments 603b, 603c apply - the compression map is referred to 508. If the read data is compressed, multiplexer 618 selects the output of the decompression logic circuitry 616b, 616c (noting that the input to the decompression logic circuitry is along a data path output of the memory interface(s) 605b, 605c). If the read data is not compressed, the multiplexer selects a data path that flows from the memory interface 605b, 605c without the decompression logic circuitry being invoked along the way.

[00104] **Figure 6c** shows an embodiment that includes compression logic circuitry 616d as well as decompression logic circuitry 616c. The compression logic circuitry 616d is shown coupled to a memory request queue 604c. As such, any compressible companion lines worth of information that are observed (or referenced) in the memory request queue 604c can be compressed together before being written into system memory. Line 617 indicates that, additionally, any cache line's worth of information waiting to be written into system memory may be compressed with its companion even if its companion is located (or referenced) in a request queue or a response queue. Additionally or in the

alternative, compression logic circuitry may be coupled to the request/response queues 602b, 603c.

Multi-Processor Systems and Cache Coherency Rules

Multiprocessor Systems

[00105] **Figure 7** illustrates a multiprocessor computing system architecture that is geared toward promoting the compression of multiple cache lines worth of data into a single cache line. A memory controller capable of sending/receiving compressed cache lines' worth of data and that maintains compression map information (e.g., as discussed above in **Figure 4a**) may be integrated into the particular multiprocessor architecture disclosed in **Figure 7**. According to the approach of **Figure 7**, compression/decompression logic 707₁, 707₂ is located within a computing sub system's 711₁, 711₂ "hub" 710₁, 710₂ so that compression may be applied during different types of transfers of instruction/data elements within the computing system. Exemplary computing system transfers that employ compression are illustrated further below with respect to **Figures 8a** through **8c**.

[00106] A processing core (e.g., processing core 701₁) is viewed as the logic circuitry devoted to the execution of instructions; and, a processor (e.g., processor 716₁) is viewed as circuitry containing a processing core (e.g., a processing core 701₁), a compression/decompression hub (e.g., hub 710₁) and a cache controller (e.g., cache controller 705₁). Irregardless of perspective, a

processor or central processing unit (CPU) typically includes a processing core and cache controller circuitry for controlling at least one cache (in those systems having a cache). The perspective of **Figure 7** is consistent with this view (noting that an additional cache controller and cache may be added for each level of cache that exists in the computing system). The perspective of **Figure 7** assumes that the compression/decompression hub is part of the processor. However, other perspectives may choose to view the hub, or any of its constituent parts as a separate component from the processor.

[00107] **Figure 8a** shows that: 1) an uncompressed cache line's worth of information from a processing core (e.g., processing core 701₁) may be compressed 801 with its companion as it is passed to cache (e.g., cache 702₁); 2) that an uncompressed cache line's worth of information from a processing core (e.g., processing core 701₁) or cache (e.g., cache 702₁) may be compressed 802 as it is passed to a system memory (e.g., system memory 704₁); and, 3) that an uncompressed cache line's worth of information from a processing core (e.g., processing core 701₁) or cache (e.g., cache 702₁) may be compressed 803 as it passed to another computing sub-system (e.g., computing system 711₂).

[00108] As an example of the second case above, if an uncompressed cache line's worth of information is evicted from cache 702₁ so as to be passed from the cache controller 705₁ through the compression/decompression hub 710₁ and then to system memory controller 703₁ (so that it can be written into system memory 704₁), on its way through compression/decompression hub 710₁, the

uncompressed evicted cache line's worth of information may be compressed 802 with its companion if the compression/decompression hub 710₁ has access to its companion and the pair of cache lines' worth of information can be compressed into a single cache line's worth of information.

[00109] As an example of third case above, if an uncompressed cache line's worth of information is evicted from cache 702₁ and is to be stored in system memory 704₂ rather than system memory 704₁ (so as to be passed from cache controller 705₁ through compression/decompression hubs 710₁ and 710₂ and then to system memory controller 703₂), a transfer from computing sub-system 711₁ to computing sub-system 711₂ takes place. Here, the evicted cache line's worth of information might be compressed at hub 710₁ or 710₂ (again, if a hub has access to the evicted cache line's companion and the two are compressible). In either case the evicted cache line's worth of information is deemed to be compressed as part of its transfer between computing sub systems 711₁ and 711₂. Thus, as depicted in **Figure 8a**, an evicted cache line's worth of information may be compressed as part of a transfer to its corresponding computing unit's system memory 802; or, as part of a transfer outside its computing sub system 803.

[00110] In the computing system observed in **Figure 7**, each compression/decompression hub includes compression/decompression logic circuitry 707₁, 707₂ and a buffer 712₁, 712₂ that temporarily queues the cache lines' worth of information it has recently received in hope that a pair of uncompressed and compressible companions are resident at the same

compression/decompression hub, at the same time, so as to enable the companions to be compressed.

[00111] **Figure 8b** demonstrates that, similarly, an uncompressed cache line's worth of information that was read from a system memory may be compressed with its companion as it passes through a compression/decompression hub while on its way to its resident computing sub system's cache 804 or to another computing sub system 805. For example, an uncompressed cache line's worth of information read from system memory 704₁ might be compressed with its companion by compression/decompression hub 710₁ en route to cache 702₂; or, by either of compression/decompression hubs 710₁, 710₂ as part of its transfer from computing sub-system 711₁ to computing sub-system 711₂.

[00112] **Figure 8c** demonstrates that an uncompressed cache line's worth of information sent from an I/O component (e.g., a magnetic disk storage unit or network interface) can be compressed with its companion (noting that a specific system memory block should be associated with the information coming from the I/O unit so that its companion can be recognized) en route to a cache 806, system memory 807 or other computing sub-system 808. For example, a cache line's worth of information received at I/O interface 708₁ may be compressed with its corresponding companion by compression/decompression hub 710₁ en route to cache 702₁, system memory 704₁ or even computing sub-system 711₂.

[00113] Computing sub-systems 711₁, 711₂ are also communicatively coupled together. Here, the structure 709 used for coupling the sub-systems to one another may be implemented as a bi-directional point-to-point link, a pair of

oppositely directed unidirectional point-to-point links, or a bus. More computing sub-systems can be added simply by adding another communicative structure between each pair of computing sub-systems. An I/O interface such as any of I/O interfaces 708₁ and 708₂ may be implemented with a point-to-point link or a bus coupled to an I/O component. More than one I/O component may be coupled to any single compression/decompression hub.

[00114] Over the course of system operation, compressed and uncompressed cache lines' worth of information are passed over interfaces to those system components that are capable of receiving both types of information (i.e., compressed information and uncompressed information); and, only uncompressed cache line's worth of information are passed to those system components that are only capable of receiving uncompressed information. For example, consider an implementation where the processing cores 701₁, 701₂ and the I/O components (not shown in **Figure 7**) are only capable of receiving (and working with) uncompressed cache lines' worth of information.

[00115] In this case, only uncompressed cache lines' worth of information would be passed to the processing cores and the I/O units. Assuming the caches 702_{1,2}, cache controllers 705_{1,2}, system memory controllers 703_{1,2} and system memories 704_{1,2} are each capable of receiving uncompressed and compressed cache line's worth of information, uncompressed and compressed cache line's worth of information would be passed to these devices. Here, for each cache line's worth of information that is at least sent to a system component capable of

receiving either compressed or uncompressed information some type of identifier would be associated with it to identify whether or not it is compressed.

[00116] According to a basic embodiment, any cache line's worth of information that is in transit in the system is appended with header information that identifies whether the cache line's worth of information is "compressed" or "uncompressed" (e.g., a bit in a header that is set active if compressed or set inactive if uncompressed, or a special encoding of a subset of the header information).

While the cache line's worth of information is within a basic component of the system (e.g., a cache, a cache controller, a memory controller, a compression/decompression hub, etc.) it is apt to be buffered/queued or stored. While the cache line's worth of information is being buffered/queued or stored, a bit or encoded value may be set so as to identify whether or not the cache line's worth of information is "compressed" or "uncompressed".

[00117] The bit or encoded value may be part of a header that is appended to the cache line's worth of information while queued or stored; or, may be referenced to its storage or buffer/queue position. In one embodiment, the bit or encoded value is located in an Error Correction Coding (ECC) field (which typically "follow" groups of information (e.g., as part of header information)). In another application of an ECC field, the ECC fields of system memory are used as a "backup" to the compression map. Here, in various embodiments, ECC information is stored with all data in system memory.

[00118] Perhaps even more advantageous, certain references to the compression can be completely avoided (e.g., such as reference 508 in

methodology 553) by recording within an ECC field that is applicable to (in part or in whole) a particular cache line's worth of information that is stored in system memory whether or not the particular cache line's worth of information is "compressed" or "uncompressed". Here, simply by reading the information and its associated ECC information from memory, the state (compressed or not compressed) of the cache line's worth of information is available. This is especially advantageous when there has been a miss to the compression map cache that has not yet been serviced by the time the corresponding data read fetch has completed, in that any further delay to service the compression map cache miss is obviated.

[00119] **Figure 9** illustrates an embodiment of a "data path" design for a compression/decompression hub 910. The embodiment of **Figure 9** can be viewed as being broken down into those interfaces that communicate with system components that do not send/receive compressed information 901 and those interfaces that communicate with system components that are capable of sending/receiving both compressed and uncompressed information 902. For example, consistent with the example discussed immediately above (an implementation where the processing cores 701₁, 701₂ and the I/O components (not shown in **Figure 7**) are only capable of receiving (and working with) uncompressed cache lines' worth of information) and where the hub design of **Figure 9** corresponds to hub 710₁ of **Figure 7**, interfaces 901 would be used to communicate to processing core 701₁ and the I/O component(s) to which computing sub-system 711₁ is coupled to (e.g., link(s) or bus 914 corresponds to

link(s) or bus 714 and link(s) or bus 908 corresponds to link(s) or bus 708₁).

Likewise, interfaces 902 would be used to communicate to: 1) the computing sub-system 711₂; 2) cache controller 705₁; and, 3) system memory controller 703₁ (e.g., link(s) or bus 909 corresponds to link(s) or bus 709; link(s) or bus 913 corresponds to link(s) or bus 713; and link(s) or bus 917 corresponds to link(s) or bus 717).

[00120] According to the data path design of **Figure 9**, a received uncompressed cache line's worth of information that is sent from any system component source and is to be sent to a system component target that is capable of receiving compressed cache lines' worth of information is entered into buffer 912 (which can be viewed for the sake of example as corresponding to buffer 712 of **Figure 7**). If the companion to the uncompressed cache line's worth of information also happens to reside in buffer 912 during the pendency of the uncompressed cache line's worth of information within buffer 912, and if the companions are compressible with one another, they are compressed together by compression logic circuitry 907a.

[00121] If the companion is not observed in buffer 912 during the pendency of the uncompressed cache line's worth of information within buffer 912, the compression logic circuitry is by passed with data path 920. Depending on whether or not compression occurs determines which channel of multiplexer 905 is selected. Upon being forwarded from multiplexer 905, the compressed or uncompressed cache line's worth of information is forwarded to interfaces 902 and sent to the appropriate target system component. Note that, again, in the

process described above the originally received uncompressed cache line's worth of information may be received from a source that does or does not send compressed information. In the former case, data path 922a is used. In the later case data path 921 is used. As the target system component is capable of receiving compressed or uncompressed information, interfaces 902 are used as an emission point from the hub 910.

[00122] Also, according to the data path design of **Figure 9**, a received cache line's worth of information that is sent from a system component source that is capable of sending compressed (i.e., is received at interfaces 902) and is to be sent to a system component target that is not capable of receiving compressed cache lines' worth of information (i.e., is to be sent from interfaces 901) is checked to see if it is compressed or not. If it is compressed, the compressed cache line's worth of information is sent to decompression logic circuitry 907b along data path 922b and is decompressed.

[00123] If it is not compressed, the uncompressed cache line's worth of information bypasses the decompression logic circuitry 907b by being passed along data path 922c. If decompression occurs, the resulting uncompressed cache lines' worth of information from decompression logic circuitry 907b are forwarded to the target system component through interfaces 901. If decompression does not occur, the originally received uncompressed cache line's worth of information is forwarded to the target system component through interfaces 901. Depending on whether or not compression occurs determines which channel of multiplexer 906 is selected.

[00124] Referring back to **Figure 7**, depending on implementation, the cache controllers 705 and/or memory controllers 703 may or may not also include their own logic circuitry for performing compression and decompression. For those situations where the target system component for a cache line's worth of information is a cache controller or system memory controller that does not include compression circuitry, a hub's compression circuitry logic circuitry (such as compression logic circuitry 907a of **Figure 9**) can be used to compress information before it is stored in cache or system memory, respectively.

[00125] In this case, in order to enhance the likelihood of storing a cache line's worth of information in a compressed format in cache or system memory, a copy of a cache line's worth of information that has been sent to a cache controller or system memory controller for purposes of being written into the corresponding cache or system memory is "held" in buffer 912 for a period of time. As such, the companion to the cache line's worth of information can be compressed with the copy of the cache's line worth of information by compression logic circuitry 907a and the compressed result stored in cache or system memory if: 1) the companion appears in the buffer 912 before the copy of the cache line's worth of information is extinguished from the buffer 912; 2) the companion is being sent to the same cache or system memory component for the purposes of being written therein; and, 3) the companion is compressible with the cache line's worth of information.

[00126] Moreover, for those situations where the source system component for an uncompressed cache line's worth of information is a cache controller or

system memory controller that does not include compression circuitry (i.e., an uncompressed cache line's worth of information is read from a cache or system memory whose corresponding controller does not contain compression circuitry), the companion to the uncompressed cache line's worth of information may also be read from the source cache or system memory and both the read uncompressed cache line's worth of information and its read companion are stored in buffer 912. If they are compressible they are subsequently compressed by compression logic circuitry 907a. If they are not compressible, the uncompressed cache line's worth of information bypasses compression logic circuitry 907a and the companion is extinguished from the buffer 912. In further embodiment, the reading of the companion is not performed and the possibility of compression is not engaged if the target for the uncompressed cache line's worth of information is a destination that does not accept compressed information.

[00127] Even though the compression methodologies described just above ("holding of copy in buffer", "reading of companion") can be used in systems whose cache controller(s) and/or memory controller(s) do not contain compression circuitry, it is conceivable that these same methodologies could be used in systems whose cache controller(s) and/or memory controller(s) do include compression circuitry. Assuming that little performance advantage is gained by so doing, however, these methodologies may be "dropped" from the capabilities of the system if the cache controller(s) and/or memory controller(s) do in fact contain compression circuitry. Moreover, in various embodiments, the cache controller(s) and/or memory controller(s) that are capable of storing

compressed information in their corresponding cache(s) and/or system memory component(s) are designed so as to include decompression circuitry and supporting circuitry so as to deal with a situation where a “write” of new information to a compressed cache line causes the compressed cache line to be no longer compressible.

[00128] With respect to system memory controller implementations, in various multi-processor system embodiments, the system memory controller components are non-traditional memory controllers that store compression map information in system memory such as those embodiments discussed above with respect to **Figures 4a** through **6c**. However, the manner in which the compression map is implemented may take on various forms. For example, according to a first approach (referred to as a “centralized approach”) the compression map for an entire system is viewed as a single body of information; while, in a second approach (referred to as “distributed”) the memory map is viewed as discrete pieces that collectively correspond to the compression map for an entire system. In multiprocessor, centralized compression map environments, the processors’ own caches can be used to store the “cached portion” of the compression map (as such, the compression map cache may be eliminated from the memory controller(s) altogether if desired).

[00129] Moreover, prior art multiprocessor cache coherency techniques (such as rules that define a particular MESI state for each cache line and inter-processor/bus agent communication requirements) can be used to manage the content of the cached portion of the compression map as well as communicate

changes made thereto. Typically, there is a memory controller and its own associated system memory resources for each processor in a multiprocessor system. As such, the compression map information in a multiprocessor, centralized approach may be stored across the different system memory resources in the system; and, the compression map information should be expanded from the embodiments 412a, 412b of **Figure 4b** to further identify which particular system memory within the system is associated with a particular block of system memory to which the compression map identifies as being compressed or non compressed.

[00130] In multiprocessor, distributed compression map environments, each memory controller in the system can be configured to maintain a portion of the compression map. Here, each controller's compression map identifies the compression/decompression status of the various blocks of the system memory resources that it is responsible for controlling the access to. Thus, compression map embodiments 412a, 412b are suitable for a distributed approach. Moreover, if the cached portion of the compression map is not stored in the processors' caches, the memory controllers may be designed so as to include a compression map cache.

[00131] **Figure 10** shows an embodiment of a map of the address bits used to reference a cache line's worth of information from a cache that is capable of storing both compressed and uncompressed cache lines' worth of information. The low-order bits 1000 determine the byte offset within the cache line. The next-most low-order bits 1001 are used either: 1) to select among companions

that have been compressed together, in the case of a compressed cache line entry; or 2) as part of the cache tag, in the case of a non compressed cache line entry. The next-most low-order bits 1002 are set bits, that select the set in the cache. Finally, the high-order bits 1003 in the address specify the tag bits.

[00132] A snoop to any structure that holds compressed information treats bits 1001 just above the address bits 1000 that determine the offset within the cache as don't cares, thus matching on either the high or low half of a compressed macro block. For example, if the line size is 74 bytes (thus bits 0-7 are used for offset 1000), and 2:1 compression is used, bit 9 (which **Figure 10** shows as the select bit 1001) selects a particular cache line with compressed information. A comparison between the desired address and the stored block would ignore bit 9. If 4:1 compression were used instead, bits 9 and 10 would be ignored.

Cache Coherency Rules

[00133] Irrespective of the precise type of multiprocessor architecture environment employed cache coherency rules should be employed. When multiple processors have the ability change the data elements that they share, it is important to keep track of the most "up-to-date" data – otherwise a processor may accidentally operate with "stale" information. By implementing cache coherency rules, the processors within a multiprocessor environment cooperatively inform each other of the status of the data elements within a computing system; and, react to them appropriately. **Figure 7** shows cache coherency rule logic 715₁, 715₂ associated with each

compression/decompression hub 710₁, 710₂ that includes logic circuitry (e.g., circuitry that includes logic gates) for implementing cache coherency rules such as those described below. Whether or not the cache coherency rule logic is part of a compression hub may vary from embodiment to embodiment. The cache coherency rule logic should be coupled to its processor's corresponding cache controller in some fashion (if not an external interface of the corresponding processor as well (e.g., a compression/decompression hub)) so that the applicable cache coherency rules can be effected.

[00134] A well known set of states for supporting cache coherency, referred to as MESI (Modified, Exclusive, Shared, Invalid) describe four possible states that may be applied to the information within any cache line. The Modified state (M) means that the content of the cache line is different than the content of its corresponding block in system memory. The Exclusive state (E) means that the content for the cache line is only available from a single cache in the system. The Shared state (S) means that the content of the cache line may exist in other caches in the system. The content of a cache line labeled as Invalid (I) means that its content is potentially stale and should be deemed as not available.

[00135] Generally, as processors operate, they change the content of their corresponding cache lines; which, in turn, may cause the state of this content to change. Cache coherency rules therefore not only involve the assignment of specific states, but also, define state transitions that define how a cache line state is to change depending upon the activity that is performed with its content. **Figures 11 through 14** review some cache coherency state transitions that may

be employed in a multiprocessor computing system that embraces compression at the cache line level. Here, it is to be understood that a compressed cache line's worth of information may be given any of the M, E, S or I states.

[00136] Figure 11 illustrates a state transition table embodiment for those situations where a compressed cache line's worth of information is written to with new information; and, the cache line's worth of information is able to maintain its compressed state (i.e., the newly written information did not cause the content to be deemed uncompressible). According to the state transition table of **Figure 11**, the changing of a compressed cache line's content that still allows the content to remain compressed causes the compressed cache line to reach the Modified state (M) regardless of which state it was previously within (e.g., M, E, S or I).

[00137] Figure 12 illustrates an embodiment of a state transition table for those situations where a compressed cache line's worth of information is written to with new information; but, the cache is not able to maintain its compressed state (i.e., the newly written information causes the content to be split into its constituent companion cache line's worth of information). According to the state transition table of **Figure 12**, again, irrespective of the state of the original compressed cache line's worth of information (i.e., X = "don't care"), the two companions that are produced are initially placed into the Modified (M) state.

[00138] Figure 13 illustrates an embodiment of a state transition table that is employed when a pair of single companion lines are compressed together to form a single compressed cache line. Here, it is assumed that the first and

second companions already exist in cache and that a change is to be written to the content of one of the cache lines that causes the pair of companions to be deemed compressible. Here, if the state of the companion to the contents of the cache line to be written to is in the M or S state the compressed cache line's worth of information that is produced as a result is placed into the M state. Note that if the state of the companion prior to the write is the S state, the compression should not take place unless a successful request for ownership (RFO) occurs with respect to the S state companion. That is, whatever other entity is using the companion (e.g., so as to cause its being in the S) state allows to "give up" its ability to change it.

[00139] Figure 14 illustrates another embodiment of a state transition table that is employed when a pair of companion lines are compressed together to form a single compressed cache line's worth of information. Here, it is assumed that: 1) a cache line's worth of information is fetched from an external resource (e.g., a system memory or a cache from another computing unit) whose companion is already resident in the cache for whom the fetched cache line's worth of information was fetched; and, 2) that the pair of companions are compressible. According to the state transition table of **Figure 14**, irregardless of what state the companion is in prior to the read (i.e., X = "don't care"), the state of the compressed cache line's worth of information is M. Note also that for a companion in the M state, compression should not take place unless a successful RFO for it occurs.

[00140] **Figure 15** shows a cache coherency rule that arises from a Request For Ownership (RFO). An RFO is a request by a processor for a specific item of information that lets other processors (and/or other components within the computing system) know that the requesting processor intends to use the specific item for its own exclusive use. Accordingly, in the methodology of **Figure 15**, a first processor requests 1501 a compressed cache line's worth of information for its own exclusive use by way of an RFO. In response, if the requested compressed cache line's worth of information is located within a cache that is associated with another (i.e., second) processor, the second processor transfers the compressed cache line's worth of information to the first processor and causes the compressed cache line's worth of information to be labeled in the invalid state (I) within its own cache 1502.

Cache Capable of Compressing/Decompressing Information

[00141] As the above described system embodiments may employ a cache controller that is capable of compressing and decompressing cache lines' worth of information, the discussion surrounding Figures 16 through 20 relate to cache controller embodiments capable of compression/decompression activity. **Figure 16A** illustrates an exemplary memory address implemented in an traditional cache. In a traditional cache, an address is divided according to tag, set and offset components. The set component is used to select one of the sets of lines. Similarly, the offset component is the low order bits of the address that are used

to select bytes within a line. **Figure 16B** illustrates one embodiment of a memory address implemented for lookup in a cache capable of working with compressed information (hereinafter a "compressed cache"). **Figure 16B** shows the implementation of a companion bit used to map companion lines of information into the same set. The companion bit is used in instances where a line of information is not compressed. Accordingly, if a line of information is not compressed, the companion bit indicates which of the two compressed companion lines of information are to be used.

[00142] In one embodiment, the window of address bits that are used for set selection is shifted to the left by one so that the companion bit lies between the set selection and byte offset bits. In this way, companion lines map to the same cache set since the companion bit and set selection bits do not overlap. The companion bit, which now is no longer part of the set selection bits, becomes part of the tag, though the actual tag size does not increase. In a traditional uncompressed cache, the companion bit is a part of the address and is used in set selection to determine whether an address hashes to an odd or even cache set.

[00143] **Figure 17** illustrates one embodiment of a tag array entry for a compressed cache. The tag array entries include the companion bit (e.g., as part of the address tag bits) and a compression bit. The compression bit causes the compressed cache tag to be one bit larger than a traditional uncompressed cache's tag. The compression bit indicates whether a line of information is compressed. Particularly, the compression bit specifies how to deal with the

companion bit. If the compression bit indicates a line of information is compressed, the companion bit is treated as a part of the offset because the line is a compressed pair. If the compression bit indicates no compression, the companion bit is considered as a part of the tag array and ignored as a part of the offset.

[00144] **Figure 18** is a block diagram illustrating one embodiment of cache controller 1804. Cache controller 1804 includes set and way selection logic 1810, byte selection logic 1820 and compression logic 1830. Set and way selection logic 1810 is used to select cache lines within a compressed cache. **Figure 19** illustrates one embodiment of set and way selection logic 1810 in a compressed cache. Referring to **Figure 19**, set and way selection logic 1810 includes tag comparison logic 1910 that receives input from a tag array to select a cache line based upon a received address. The tag comparison logic 1910 takes into account whether a cache line holds compressed data.

[00145] Because cache lines can hold a variable data size, tag comparison logic 1910 is also variable length, depending on whether a particular line is compressed or not. Therefore, the tag match takes into account the compression bit. When compressible by at least 2:1, the two sectors of each line are stored in a single physical cache line (e.g., in one way). It is important to note that this differs from traditional sectored cache designs in that different logical sectors of a given logical line may be stored simultaneously in different ways when uncompressed.

[00146] According to **Figure 18**, byte selection logic 1820 selects the addressed datum within a line. According to one embodiment, byte selection logic 1820 depends on the compression bit. **Figure 20** illustrates one embodiment of byte selection logic 1820. Byte selection logic 1820 includes a decompressor 2010 to decompress a selected cache line if necessary. An input multiplexer selects between a decompressed cache line's worth of information and an uncompressed cache line's worth of information depending upon the compression bit. In one embodiment, the range of the offset depends on whether the line of information is compressed.

[00147] If the line of information is compressed, the companion bit of the address is used as the high order bit of the offset. If the line of information is not compressed, decompressor 2010 is bypassed and the companion bit of the address is not used for the offset. The selected line is held in a buffer whose size is twice the physical line size to accommodate compressed data. Alternative embodiments may choose to use the companion bit to select which half of the decompressed word to store in a buffer whose length is the same as the physical line size. However, buffering the entire line of information is convenient for modifying and recompressing data after writes to the cache.

[00148] Compression logic 1830 may also be used to determine when a line of information is to be compressed. According to one embodiment, opportunistic compression is used to determine when a line of information is to be compressed. The above-described mechanism allows any two cache line's worth of information that map to the same set and that differ only in their

companion bit to be compressed together into one cache line. In one embodiment, the mechanism modifies the set mapping function and selects the companion bit such that it allows adjacent memory lines of information to be compressed together, which takes advantage of spatial locality.

Closing Comments

[00149] Note also that embodiments of the present description may be implemented not only within a semiconductor chip but also within machine readable media. For example, the designs discussed above may be stored upon and/or embedded within machine readable media associated with a design tool used for designing semiconductor devices. Examples include a circuit description formatted in the VHSIC Hardware Description Language (VHDL) language, Verilog language or SPICE language. Some circuit description examples include: a behavioral level description, a register transfer level (RTL) description, a gate level netlist and a transistor level netlist. Machine readable media may also include media having layout information such as a GDS-II file. Furthermore, netlist files or other machine readable media for semiconductor chip design may be used in a simulation environment to perform the methods of the teachings described above.

[00150] Thus, it is also to be understood that embodiments of this invention may be used as or to support a software program executed upon some form of processing core (such as the Central Computing unit (CPU) of a computer) or otherwise implemented or realized upon or within a machine readable medium.

A machine readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine readable medium includes read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

[00151] In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.